US 20020161861A1

(19) United States

(12) Patent Application Publication (10) Pub. No.: US 2002/0161861 A1

Greuel (43) Pub. Date: Oct. 31, 2002

(54) METHOD AND APPARATUS FOR CONFIGURABLE DATA COLLECTION ON A COMPUTER NETWORK

(76) Inventor: James R. Greuel, Fort Collins, CO (US)

Correspondence Address:
HEWLETT-PACKARD COMPANY
Intellectual Property Administration
P.O. Box 272400
Fort Collins, CO 80527-2400 (US)

(21) Appl. No.: 09/794,376

(22) Filed: Feb. 27, 2001

Publication Classification

(51) Int. Cl.[7] G06F 15/177

(52) U.S. Cl. 709/220; 713/2

(57) ABSTRACT

The invention facilitates improved configuration of computer network data collection. In one respect, the invention is a method for determining a network data collection configuration for an entity utilizing a computer network. The method determines which nodes in the network are associated with the entity, reads a view definition associated with the entity and determines which network variables are components of the health score, selects from the nodes associated with the entity those nodes that are components of the view, and produces a configuration file that specifies the network variables for each corresponding selected node. The data collection preferably utilizes the SNMP protocol. The determining step preferably queries a database that stores entity-node associations. The view definition is preferably an XML file and refers to all nodes of a particular device type, so that the selecting step can filter network nodes by device type. In a preferred form, the network is the Internet, the method is performed by an Internet service provider, and the entity is a customer of the Internet service provider. Optionally, the method reads and utilizes a template for the configuration file. In another respect, the invention is an apparatus. The apparatus comprises a first database that stores entity-node associations, a view definition associated with an entity, a second database that stores node device type information, and a network data collection configuration generation tool, connected to the above. The network data collection configuration generation tool generates a network data collection configuration file customized to the view definition.

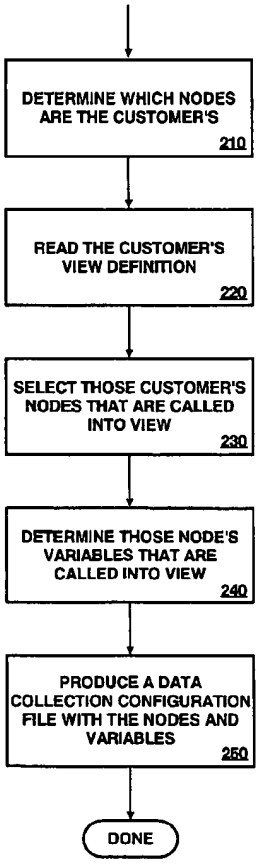

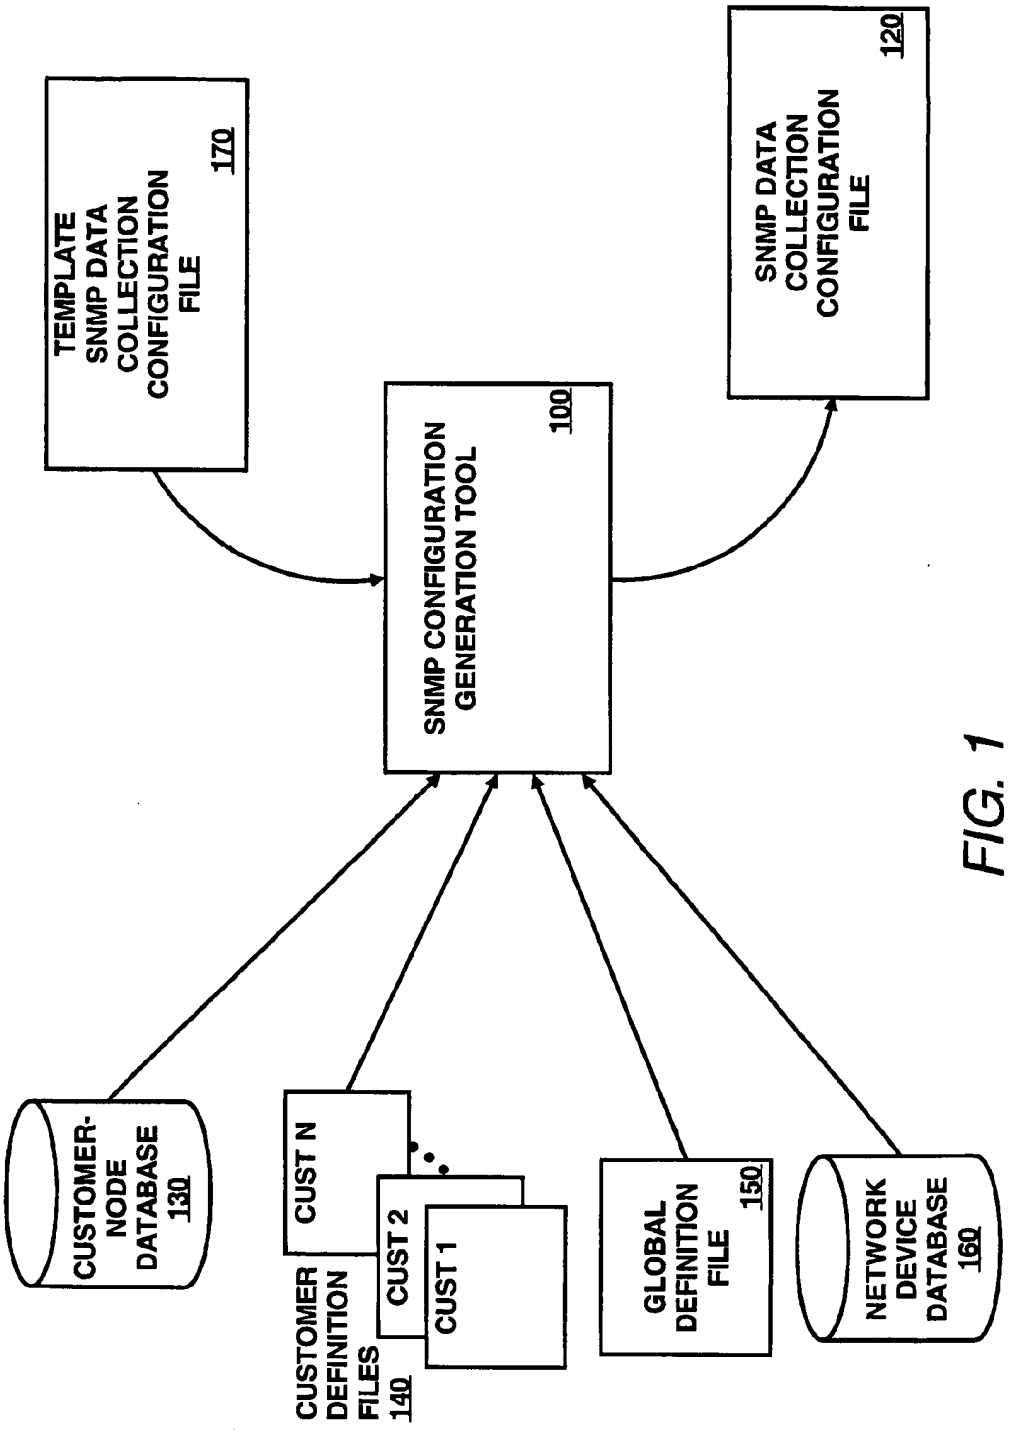
TEMPLATE SNMP DATA COLLECTION CONFIGURATION FILE
170
SNMP DATA COLLECTION CONFIGURATION FILE
120
SNMP CONFIGURATION GENERATION TOOL
100
CUSTOMER-NODE DATABASE
130
CUST N
CUST 2
CUST 1
CUSTOMER DEFINITION FILES
140
GLOBAL DEFINITION FILE
150
NETWORK DEVICE DATABASE
160

FIG. 1

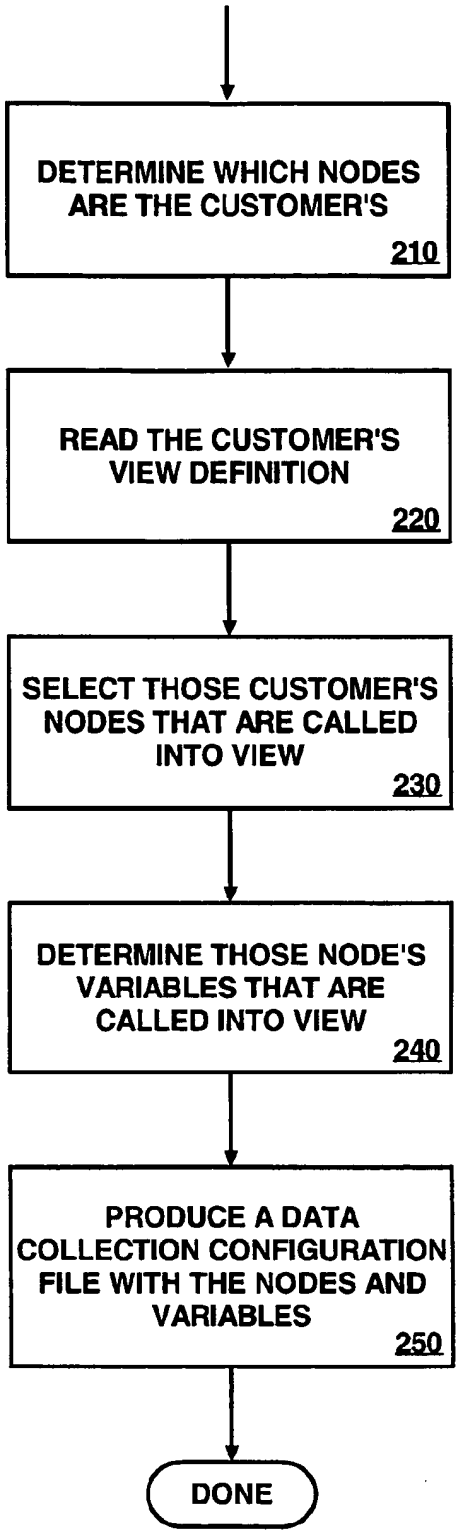
DETERMINE WHICH NODES ARE THE CUSTOMER'S
210
READ THE CUSTOMER'S VIEW DEFINITION
220
SELECT THOSE CUSTOMER'S NODES THAT ARE CALLED INTO VIEW
230
DETERMINE THOSE NODE'S VARIABLES THAT ARE CALLED INTO VIEW
240
PRODUCE A DATA COLLECTION CONFIGURATION FILE WITH THE NODES AND VARIABLES
250
DONE

FIG. 2

METHOD AND APPARATUS FOR CONFIGURABLE DATA COLLECTION ON A COMPUTER NETWORK

FIELD OF THE INVENTION

[0001] This invention relates generally to computer networks and more particularly to computer network monitoring.

BACKGROUND OF THE INVENTION

[0002] The Simple Network Management Protocol (SNMP) originated as a means for managing TCP/IP (Transmission Control Protocol/Internet Protocol) and Ethernet networks. Today, SNMP is a generally accepted standard for monitoring and control of heterogeneous networks. Using SNMP, network administrators can address queries and commands to network nodes and devices. SNMP monitors network performance and status; controls operational parameters; and reports, analyzes and isolates faults. The protocol accomplishes these functions by transporting management information between "managers" and "agents."

[0003] An agent is a component housed within a managed network device such as a router, gateway, or server. Each agent stores management data and responds to the manager's requests for this data, and may send a "trap," a special unsolicited SNMP message, to the manager after sensing a prespecified condition. A manager is a component housed within a network management station. The manager queries and controls agents using various SNMP commands, which are transported as Protocol Data Units (PDUs). Five SNMP commands are defined in SNMP version 1: GetRequest, GetNextRequest, SetRequest, GetResponse and Trap. Agents inspect and retrieve the management data after receiving either a GetRequest or a GetNextRequest PDU from a manager. Managers use GetRequest for retrieving single values of the managed objects. The GetNextRequest is issued by the manager to begin a primitive block transfer and the agent returns the selected data with a GetResponse PDU. Managers use SetRequest commands for instructing agents to alter variables.

[0004] SNMP utilizes a virtual information store (e.g., an object database) that is referred to as Management Information Base (MIB). The MIB store is accessible to agents and manipulated via SNMP for network management. A MIB data structure defines a device's observable (e.g., discoverable or collectible) variables and controllable parameters. A router MIB, for example, may contain fields for CPU utilization, up/down status for each interface, error rates on interfaces, congestion metrics (e.g., buffer levels, latency or packet discard rates) and the like.

[0005] The manager is charged with, among other things, monitoring network performance and status, controlling operational parameters, and reporting, analyzing and isolating faults in its managed domain. A good example of an SNMP network management station (and hence a manager) is Hewlett-Packard's Network Node Manager (NNM) product.

[0006] Another protocol for network performance monitoring and troubleshooting is ICMP (Internet Control Message Protocol). The ICMP protocol supports ping or echo messages, which are round-trip messages to a particular addressed network device and then back to the originator. By issuing a ping to a network device, a manager can determine whether the network device is online or offline (i.e., up or down) on the basis of whether the ping message is returned to the manager. Because the ICMP protocol or other ping messages are universally supported, the manager can in this way determine the most important piece of status information (i.e., up/down status) for network devices that do not support the SNMP protocol.

[0007] A manager can be directed to collect specific network data via SNMP or ICMP. In other words, a manager can be directed to issue certain GetRequest or GetNextRequest commands to a particular agent. Typically, these directions are contained in a data collection configuration file, which specifies the nodes and MIB variables to request as well as when to make the request.

[0008] Data collection configuration files can be prepared manually, but manual preparation is disadvantageous in several respects. Manual configuration is time consuming and error prone. The disadvantages are compound when the data collector is a network provider, who provides networks and/or network services to its customers, such as a on-line service provider (OSP), Internet service provider (ISP) or company's own IT (information technology) group (whose "customers" are other groups of the same company). Any time a new customer is added, a customer removed, service expanded, or performance metric changed, a new or different data collection configuration is likely needed. It is unlikely that a network administrator would always remember to perform the configuration change—much less always perform the change correctly.

[0009] Automatic data collection configuration is better. A version of Hewlett-Packard's NNM provided automatic SNMP data collector configuration. Though providing some improvement over purely manual configuration, the automatic configuration feature of this product was suboptimal. In particular, automatic configuration was limited to a fixed set of SNMP metrics and did not support a customer model or unique or flexible sets of SNMP variables. In addition, its support of device type filtering was very limited. That is, one could not conveniently configure data collection from all of particular type of device (e.g., all servers or all routers). Furthermore, one could not conveniently configure data collection based on a customer-to-node mapping.

SUMMARY OF THE INVENTION

[0010] The invention facilitates improved configuration of computer network data collection.

[0011] In one respect, the invention is a method for determining a network data collection configuration for an entity utilizing a computer network. The method determines which nodes in the network are associated with the entity, reads a view definition associated with the entity and determines which network variables are components of the customer view, selects from the nodes associated with the entity those nodes that are components of the view, and produces a configuration file that specifies the network variables for each corresponding selected node. The data collection preferably utilizes the SNMP protocol. The determining step preferably queries a database that stores entity-node associations. The view definition is preferably an XML (extensible markup language) file and refers to all nodes of

a particular device type, so that the selecting step can filter network nodes by device type. In a preferred form, the network is the Internet, the method is performed by an Internet service provider, and the entity is a customer of the Internet service provider. Optionally, the method reads and utilizes a template for the configuration file.

[0012] In another respect, the invention is computer readable media on which is embedded a program that performs the above method.

[0013] In yet another respect, the invention is an apparatus. The apparatus comprises a first database that stores entity-node associations, a view definition associated with an entity, a second database that stores node device type information, and a network data collection configuration generation tool, connected to the above. The network data collection configuration generation tool generates a network data collection configuration file customized to the view definition. The data collection preferably utilizes the SNMP protocol. Optionally, the view definition associated with the entity comprises an expression defined in a second view definition. Optionally, the apparatus also includes a template network data collection configuration file.

[0014] In comparison to known prior art, certain embodiments of the invention are capable of achieving certain advantages, including some or all of the following: (1) data collection configuration can be customized more easily; (2) customer models are better supported; and (3) improved device type filtering is provided. Those skilled in the art will appreciate these and other advantages and benefits of various embodiments of the invention upon reading the following detailed description of a preferred embodiment with reference to the below-listed drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

[0015] **FIG. 1** is a block diagram of an apparatus and its environment, according to an embodiment of the invention; and

[0016] **FIG. 2** is a flowchart of a method, according to an embodiment of the invention.

DETAILED DESCRIPTION OF A PREFERRED EMBODIMENT

[0017] An embodiment of the invention is explained below in the context of a customer model, because the embodiment is naturally suited to the customer model. According to this model, a service provider (e.g., OSP, ISP or company IT group) provides a computer network or part thereof and/or network services (e.g., Internet access or application hosting) for use by a customer of the service provider. The customer may be, for example, a web site host. Typically, a service provider has multiple customers. Each customer is understandably curious about its network or piece of the network. For example, each customer would like to have assurances that the service provider is providing adequate resources. Likely, each customer is curious about different aspects of the network. Although the embodiment of the invention is described with reference to "customers," the "customer" may be any entity.

[0018] **FIG. 1** is a block diagram of an SNMP configuration generation tool **100** and its environment **110**, according to an embodiment of the invention. The environment **110** includes several files and databases that are the inputs and output of the SNMP configuration generation tool **100**. The final output of the SNMP configuration generation tool **100** is an SNMP data collection configuration file **120**. The SNMP data collection configuration file **120** is input to an SNMP data collector (not shown). The SNMP configuration generation tool **100** produces the SNMP data collection configuration file **120** on the basis of input from a customer-node database **130**, customer definition files **140**, a global definition file **150**, a network device database **160** and a template SNMP data collection configuration file **170**.

[0019] The customer-node database **130** is a database that lists, for each customer, which nodes in the network are associated with that particular customer. As an example, the customer-node database **130** may contain the following information for a given customer:

[0020] Customer WingNuts-R-Us

[0021] Nodes

[0022] c8540csr.cnd.hp.com

[0023] lcntp550.cnd.hp.com

[0024] theforce.cnd.hp.com

[0025] According to the above example entries, a customer named "WingNuts-R-Us" is associated with three network nodes. The nodes may be servers, routers or other types of network devices, as specified in the network device database **160**, as explained below. What the customer wishes to observe about these nodes is specified in the customer definition files **140**, next to be described.

[0026] The customer definition files **140** are files that contain high-level definitions of what aspects of the computer network the customer desires to view, such as, performance metrics. The customer definition files **140** are also referred to as view definition files. For example, the customer definition files **140** can contain health score definitions, such as the ones described in copending and commonly assigned U.S. patent application Ser. No. 09/_____ (attorney docket No. 10006622-1), entitled "Method and Apparatus for Customizably Calculating and Displaying Health of a Computer Network," which is hereby incorporated by reference. An illustrative definition found in the customer definition file **140** is a definition for router health as follows:

```
<!-- Definition for a "Router Health" gauge. -->
<Summary title="Router Health">
    <Component weight="1" href="#IfHealth"/>
    <Component weight="1" href="#CiscoCpuUtil"/>
    <NodeSelection title="Routers">
        <CapabilityFilter>
            <Capability value="isRouter"/>
        </CapabilityFilter>
    </NodeSelection>
</Summary>
```

[0027] This example entry defines a performance metric labeled "Router Health." As defined, Router Health is to be calculated based on two equally weighted components: interface health ("IfHealth") and CPU (central processing unit) utilization ("CiscoCpuUtil"). These components, and

others that customers can utilize, are defined in the global definition file **150**, as explained in detail below. The Router Health metric is applied to all nodes that are routers (i.e., for which "isRouter" is true), according to the network device database **160**.

[0028] As can be seen from the example above, the preferred form of a customer definition file **140** is XML (extensible markup language). Each customer has its own customer definition file; this allows each customer to define its own view in a manner uniquely suitable to that customer. A customer definition file **140** contains information to identify the customer, such as the following:

```
<SecurityFilter>
    <CustomerFilter>
        <Customer name="WingNuts-R-Us"/>
    </CustomerFilter>
</SecurityFilter>
```

[0029] The global definition file **150** contains definitions for expressions that may appear in the customer definition files. The global definition file **150** is potentially referred to by all of the customer definition files **140**. As an example, the following entry defines "CiscoCpuUtil," which is a measure of CPU utilization:

[0030] <Element id="CiscoCpuUtil" title="CPU Utilization Health"

[0031] autoConfig="yes"

[0032] href="snmp://%item%[0]/p_cisco5minavgbusy">

[0033] <!—Maps utilization percentage to health scores 0-100 —>

[0034] <!—Note that higher utilization results in a lower score.—>

[0035] <Scale lower="0" upper="60" translation="100"/>

[0036] <Scale lower="60" upper="70" translation="75"/>

[0037] <Scale lower="70" upper="80" translation="50"/>

[0038] <Scale lower="80" upper="90" translation="25"/>

[0039] <Scale lower="90" translation="0"/>

[0040] </Element>

[0041] This sample entry indicates that the data source is SNMP and, more specifically, the MIB variable p_cisco5minavgbusy (which is an alias for "cisco.local.l-system.avgBusy5"). The "[0]" indicates that the SNMP instance value is zero. The "autoconfig" attribute would, if "no," cause the SNMP configuration generation tool **100** to ignore this element. The sample entry also includes a scale that maps ranges of the raw data value of the MIB variable into scores from zero to 100. Those skilled in the art realize that other forms of mappings (e.g., mathematical formulas) are possible.

[0042] The network device database **160** contains information about devices in the network. This information includes device type information, such as the following:

[0043] c8540csr.cnd.hp.com

[0044] isRouter=true

[0045] isServer=false

[0046] isHub=false

[0047] isKeyDevice=true

[0048] isCPE=false

[0049] lcntp550.cnd.hp.com

[0050] isRouter=true

[0051] isServer=false

[0052] isHub=false

[0053] isKeyDevice=true

[0054] isCPE=true

[0055] theforce.cnd.hp.com

[0056] isRouter=false

[0057] isServer=true

[0058] isHub=false

[0059] isKeyDevice=false

[0060] isCPE=false

[0061] According to these sample entries, the devices whose domain name are c8540csr.cnd.hp.com and lcntp550.cnd.hp.com are routers and also designated as "Key Devices." The device at theforce.cnd.hp.com is a server. This information enables the SNMP configuration generation tool **100** to access all devices of a particular type easily and conveniently, as the example customer definition files **140** above demonstrates. The network device database **160** preferably also includes information about device interfaces (e.g., whether a given interface is "Key").

[0062] The template SNMP data collector configuration file **170** is the final input to the SNMP configuration generation tool **100** in the environment **110**. The template file **170** is a starting point for constructing the SNMP data collector configuration file **120**. An example entry is the temple file **170** is the following:

[0063] MIB p_cisco5minavgbusy

[0064] _NODE_900_INSTANCE_

[0065] The expressions "_NODE_" and "_INSTANCE_" are placeholders that the SNMP configuration generation tool **100** overwrites with an actual node name (or address) and instance value. The term "900" is a default polling interval in seconds.

[0066] Continuing the above series of examples to fruition results in the following entry in the SNMP data collector configuration file **120**:

[0067] MIB p_cisco5minavgbusy

[0068] lcntp550.cnd.hp.com 900 0

[0069] c8540csr.cnd.hp.com 900 0

[0070] This sample entry configures an SNMP data collector to collect values of the MIB variable p_cisco5minavgbusy from the two nodes named in the first column, with a polling frequency of 900 seconds and an instance value of 0 in both cases. Collection of these two MIB variables enables the CPU utilization component of the Router Health metric to be calculated.

[0071] Network changes can be tracked and reacted to with minimal effort in the environment **100**. If a customer is removed or a new customer is added, then the customer-node database **130** is simply changed accordingly. If devices are removed or new devices are added to the network, then the network device database **160** is changed accordingly. Expressions can be modified in the global definition file **150** without disturbing each of the customer definition files **140**. Polling intervals can also be adjusted globally by modification to the template file **170**.

[0072] Although the embodiment illustrated in **FIG. 1** has been described with reference to SNMP, the invention contemplates any type of computer network data collection protocol. One illustrative protocol other than SNMP is ICMP.

[0073] **FIG. 2** is a flowchart of a method **200** according to an embodiment of the invention. The method **200** produces the data collection configuration file **120**. The method **200** is performed by the configuration generation tool **100**. The method **200** determines (**210**) which nodes are associated with the customer. Preferably, the determination step **210** is performed by querying the customer-node database **130**. The method **200** reads (**220**) the customer's view definition. Preferably, the view definition is an XML file that defines customized network health scores that the customer wishes to view, like the customer definition files **140**. The view definition may be independent and self-contained, or it may refer to expressions defined elsewhere, such as the global definition file **150**. The method **200** next selects (**230**) those nodes associated with the customer that are called into view and determines (**240**) those node's variables that are called into view. Finally, the method **200** produces a data collection configuration file with the nodes and variables determined in the steps **230** and **240**. The production step **250** can be performed with or without the aid of a preexisting template file.

[0074] The steps of the method **200** can be performed in an order different from that illustrated in **FIG. 2**, as one skilled in the art readily appreciates. Furthermore, certain steps of the method **200** can be performed simultaneously in parallel. For example, the determining step **210** and the reading step **220** are not sequentially dependent and can be performed concurrently.

[0075] The method **200** can be performed by a computer program. That is, the SNMP configuration generation tool **100** can be a computer program. The computer program can exist in a variety of forms both active and inactive. For example, the computer program and objects can exist as software comprised of program instructions or statements in source code, object code, executable code or other formats; firmware program(s); or hardware description language (HDL) files. Any of the above can be embodied on a computer readable medium, which include storage devices and signals, in compressed or uncompressed form. Exemplary computer readable storage devices include conventional computer system RAM (random access memory), ROM (read only memory), EPROM (erasable, programmable ROM), EEPROM (electrically erasable, programmable ROM), and magnetic or optical disks or tapes. Exemplary computer readable signals, whether modulated using a carrier or not, are signals that a computer system hosting or running the computer program can be configured to access, including signals downloaded through the Internet or other networks. Concrete examples of the foregoing include distribution of executable software program(s) of the computer program on a CD ROM or via Internet download. In a sense, the Internet itself, as an abstract entity, is a computer readable medium. The same is true of computer networks in general.

[0076] What has been described and illustrated herein is a preferred embodiment of the invention along with some of its variations. The terms, descriptions and figures used herein are set forth by way of illustration only and are not meant as limitations. For example, the score calculated and output by the invention need not be a "health" score, and the score need not be a composite formed from two or more system variables, but may be a score derived from a mapping of a single system variable. Those skilled in the art will recognize that these and many other variations are possible within the spirit and scope of the invention, which is intended to be defined by the following claims—and their equivalents—in which all terms are meant in their broadest reasonable sense unless otherwise indicated.

What is claimed is:

1. A method for determining a network data collection configuration for an entity utilizing a computer network, the method comprising:

determining which nodes in the network are associated with the entity;

reading a view definition associated with the entity and determining which network variables are components of the view definition;

selecting from the nodes associated with the entity those nodes that are components of the view definition; and

producing a configuration file that specifies the network variables for each corresponding selected node.

2. The method of claim 1 wherein network data collection configuration is an SNMP data collection configuration.

3. The method of claim 1 wherein the view definition is an XML file.

4. The method of claim 1 wherein the determining step comprises:

querying a database that stores entity-node associations.

5. The method of claim 1 wherein the view definition refers to all nodes of a particular device type, and the selecting step comprises:

filtering network nodes by device type.

6. The method of claim 1 wherein the network is the Internet, the method is performed by an Internet service provider, and the entity is a customer of the Internet service provider.

7. The method of claim 1 further comprising:

reading a template for the configuration file; and

wherein the producing step comprises utilizing the template.

8. An apparatus comprising:

a first database that stores entity-node associations;

a view definition associated with an entity;

a second database that stores node device type information; and

a network data collection configuration generation tool, connected to the first database, the view definition and the second database, whereby the network data collection configuration generation tool generates a network data collection configuration file customized to the view definition.

9. The apparatus of claim 8 further comprising the network data collection configuration file.

10. The apparatus of claim 8 wherein the network data collection configuration file is an SNMP data collection configuration file.

11. The apparatus of claim 8 further comprising:

a second view definition, wherein view definition associated with the entity comprises an expression defined in the second view definition.

12. The apparatus of claim 8 further comprising:

a template network data collection configuration file, connected to the network data collection configuration generation tool.

13. The apparatus of claim 14 further comprising:

a network data collector, connected to the network data collection configuration.

14. A computer readable medium on which is embedded a program, the program performing a method for determining a network data collection configuration for an entity utilizing a computer network, the method comprising the following steps:

determining which nodes in the network are associated with the entity;

reading a view definition associated with the entity and determining which network variables are components of the view definition;

selecting from the nodes associated with the entity those nodes that are components of the view definition; and

producing a configuration file that specifies the network variables for each corresponding selected node.

15. The computer readable medium of claim 15 wherein network data collection configuration is an SNMP data collection configuration.

16. The computer readable medium of claim 15 wherein the view definition is an XML file.

17. The computer readable medium of claim 15 wherein the view refers to all nodes of a particular device type, and the selecting step comprises:

filtering network nodes by device type.

18. The computer readable medium of claim 15 wherein the network is the Internet, the method is performed by an Internet service provider, and the entity is a customer of the Internet service provider.

* * * * *